



Atmospheric Science Data Center

CERES Science Team Meeting

April 28, 2009

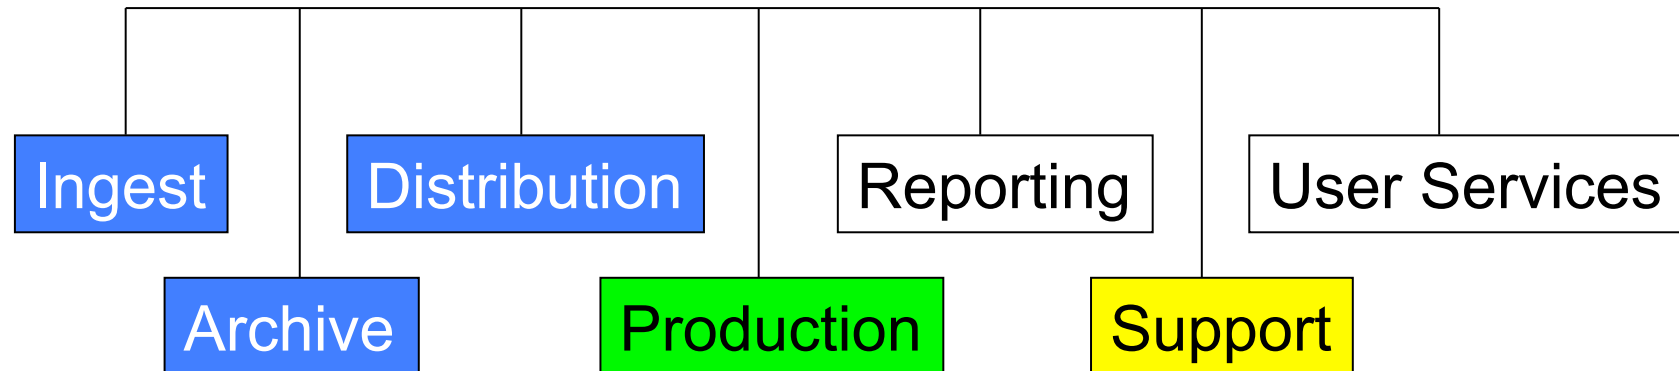
Hampton, VA

Summary



- ASDC Partners with CERES
 - Evolve CERES Data Products
 - Support Cal/Val of instruments and Data Products
 - Preserve long-term Climate Data Records
- Modernize ASDC
 - Speed access to Data Products
 - Increase capacity to re-process CERES data faster
 - Increase accessibility of Data Products
- Elements
 - ANGe – Storage Management Upgrade
 - AMI – Hardware Modernization
 - PAM – Production Automation Modernization

ASDC Breakdown



ASDC Archive Systems




- **Recent Progress**

- CERES Archive Management migrated to ANGe
- Problems with ASDC Archive Reliability resolved (StorNext)
- All Orderable CERES Products at ESDIS Clearinghouse (ECHO)
- ASDC Order Tool continues to be primary distribution path

- **Projected Progress**


- CERES Ingest functions to support FM5 Testing
- Conversion of ANGe to Disk-based archive
- Replacement of Tape Libraries with LTO-4 Tape System
- Migrate data products from tape to disk (1 year)



Production



- Science teams want faster reprocessing turn around
 - MISR looks for 5-10x improvement over previously agreed to capabilities
 - Past rates ~2x-4x
 - Current reprocessing campaign peaking >20x
 - CERES aims for aggressive rates for reprocessing (50-100x)
 - 10x reprocessing rates experienced in the past
 - Currently have over 25 instrument years of data
 - Science need to reprocess all 25 years in a year **for new version to be useful**



ASDC Production Systems




- **Recent Progress**

- Improve use of Sun Grid Engine (SGE) in production
- Hardware upgrade in progress
- Consolidation of toolkit/cereslib to single location

- **Planned Progress**


- Improve commonality of SCF and Production resources
- Continue SGE Improvements
- Complete deployment of disk-based fiberchannel archive
- Retire warlock (SGI-3800) by end of year



SCF Support to CERES




- **Tight integration of Science Computing Facility (SCF) with ASDC Resources**
 - Reliable, high speed access to Data Products by SCF
 - Consolidated set of input data, libraries, toolkits
 - Interactive processing by small groups
 - Batch processing using SGE on cluster
 - Improve persistence of user data
 - Revolving backups
 - Permit user archiving of non-production data products
- **Move SCF to B1268C**
 - Air conditioning capacity
 - Stable Electric Power (UPS)



Hardware Upgrade



- **Replace Data Product Production System**
 - Homogeneous IBM environment, designed integrated
 - ~300 Linux processing nodes acquired
 - Big-endian and Little-endian processors integrated access to disk
 - 1.2 petabytes of raw disk storage
 - High-speed connections between components
 - Fiberchannel through fast, wide switches
 - 5 years prepaid maintenance
- **Backup, Archive and disaster recovery**
 - Replace miscellaneous tape systems with LTO-4 Libraries
 - Compatibility between ASDC/SCF tape systems
 - Auto revolving backup of working files (interactive)
 - Permit users to archive non-Production data products
 - ANGe archives Production Data Products



Schedule for Capabilities



- May 1, 2009
- May 4-8, 2009
- May 14, 2009
- May 15, 2009
- May 21, 2009
- October, 2009
- December, 2009
- December 31, 2009
- SGI3800 Move to B1268c
 - Thunder/lightning replaced
- User Training: IBM, SGE
- PI to provide priority for data product migration to disk
- AMI (IBM) operational
 - SCF Disk/Processors
 - Code Delivery to ASDC
 - Data Product Processing
- FM-5 Interface test (PAN/PDR)
- Tape System Operational
- Data Migration Complete
- Warlock off maintenance




Backup Slides

AMI Initial Operational Capability




- **ASDC**
 - Production Processors available to Operators
 - Ready to accept code delivery
 - Disk System Directories created and data migration started
- **Combined**
 - Consolidated Source of cereslib, toolkit available
 - Limited NFS, AFP, SAMBA access to disks
 - Data Product File Systems
 - Data migrated from earlier disk array
 - New data products available on disk
- **SCF**
 - SCF Processors available to Authorized Users
 - Interactive blades (Power6)
 - SGE Batch processors
 - Initial increment of SCF Disk System available
 - Quotas Assigned to Groups based on Payment
 - Read-only access to ASDC Data Products




Plans for Next 1 to 2 Years




- Mature ANGe
- Finalize the setup and configuration of new processing hardware
- Prepare and reprocess MISR data
- Retire and minimize use of CERES SGI processing systems
- Tightly integrate ASDC and CERES SCF
- Automate CERES processing
- Prepare ANGe to support new data providers
- Collaborate with partners to provide for more data fusion



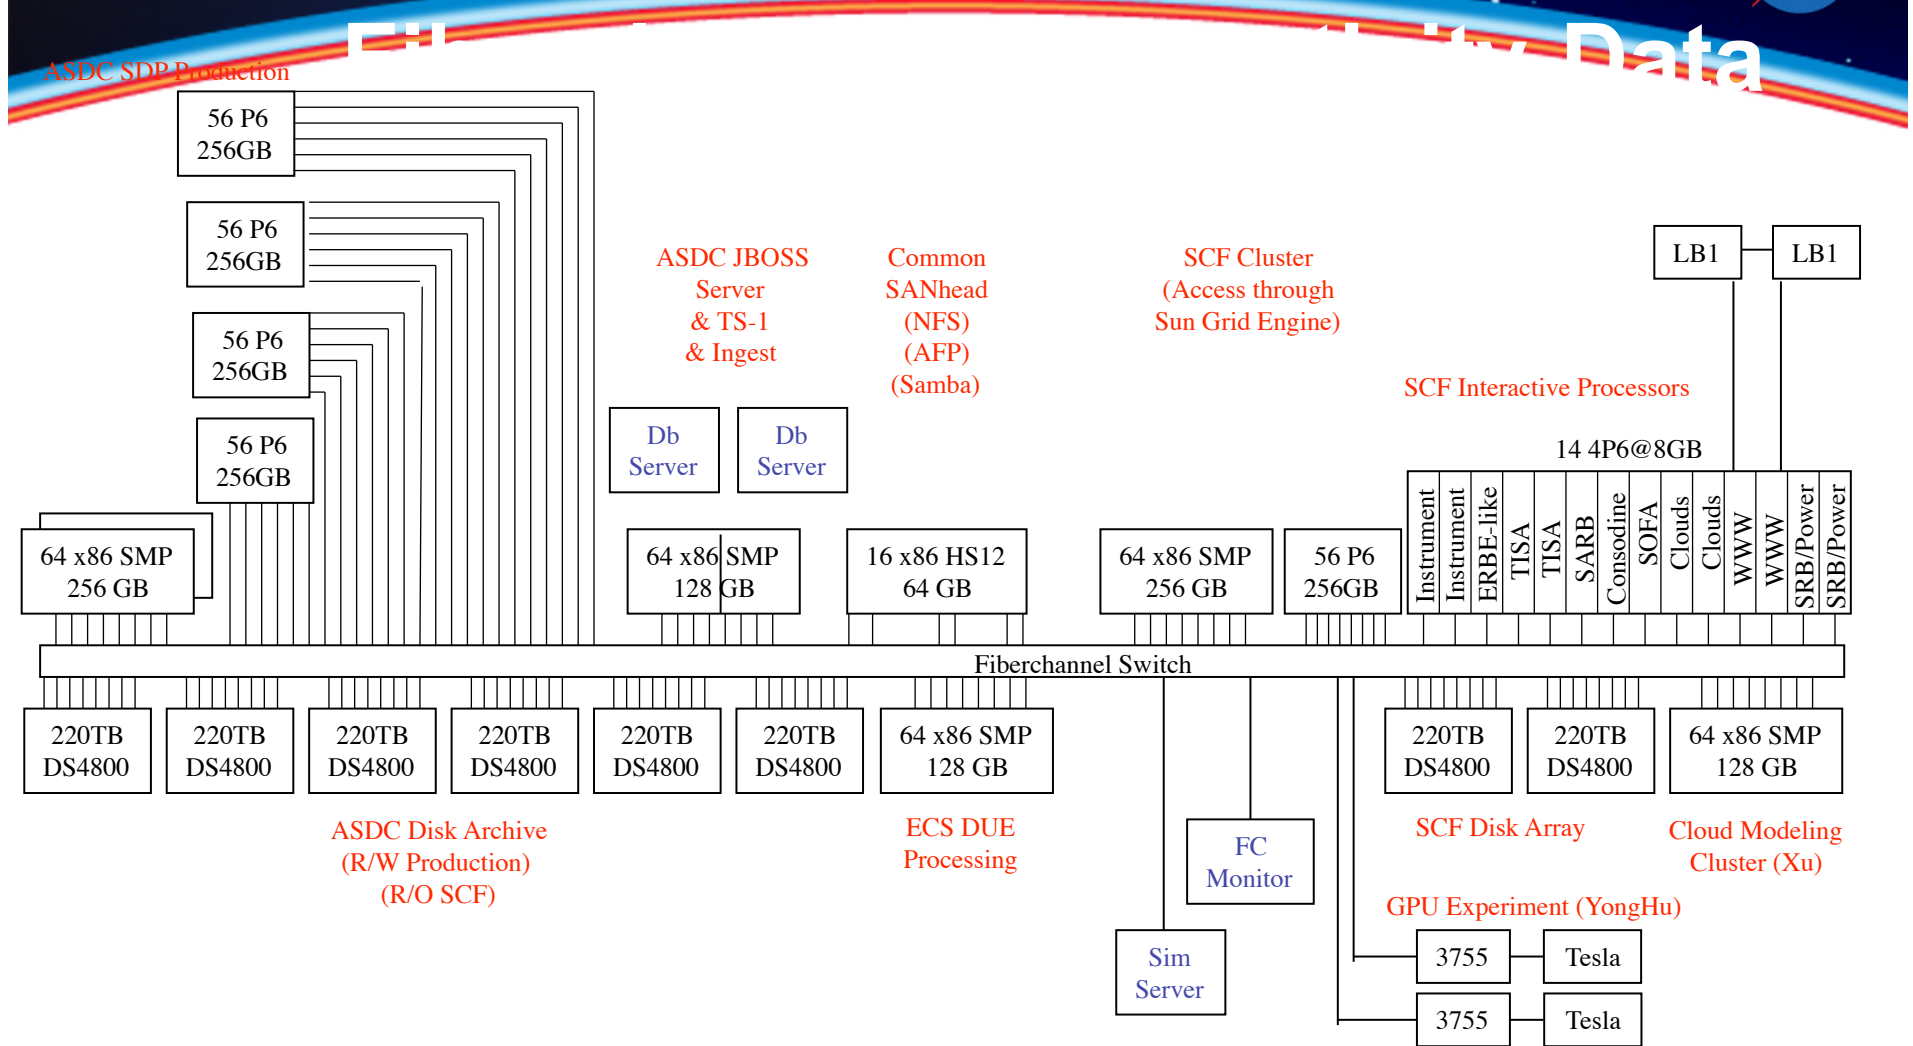
Plans for Next 3+ Years



- Continue to modernize ANGe
- Continue to modernize processing and archive environments
- Assume long term preservation of data
- Prepare and reprocess MISR data
- Explore using ANGe flexibility to support external archive requirements
- Continue to collaborate with partners to provide for more data fusion



ASDC & SCF Integrated Architecture



Basic System for ASDC



- Permits only running in production by operators, not adhoc by users
- CERES Production Processing System
 - 224 Power6 processors in LINUX (4GB per processor)
 - 128 Xeon SMP in LINUX (4GB per processor)
 - Commercial ADA for Legacy Instrument codes
 - Shared access to memory permits dynamic allocation to processor
 - IDL running on Shared License server
 - Eventually integrate legacy systems
 - 124 Power4 processors in LINUX (2GB per processor)
 - Legacy SGI3800 as interim fix
- Flashflux, SRB processing, MISR reprocessing
- JBOSS Server for ANGe
 - Production 32 Xeon SMP in LINUX (2 GB per processor)
 - Test/development environment matching production which can be re-allocated if needed to production
- Database Server
- Dropbox for ingest
- WWW cluster to deliver data and public information
- DAAC Unique Enhancements for ECS processing
 - MISR subsetting
 - Others
- Storage of Archived data
 - 900TB (usable) RAID-5 in 6 Chassis
 - 200 TB CERES Aqua+Terra
 - 150 TB MODIS Aqua+Terra
 - 150 TB CERES NPP
 - 200 TB VIIRS
 - Expansion capacity to 10 PetaBytes
- Communications
 - Internal Fiberchannel dual 4GB FC for data sharing
 - Control and Monitor via TCP/IP
 - External NFS, Samba, AFP via GB ethernet from multiple file servers
- Support
 - Installation and Checkout by IBM
 - 24x7x365 for 5 years

Basic System for SCF



• User directed Processing System

- Interactive IBM Systems
 - 56 Power6 processors assigned to subsystems (interactive) (4GB@)
- SGE-based batch IBM systems
 - 56 Power6 processors in LINUX cluster (4GB per processor)
 - 64 Xeon SMP in LINUX (4GB per processor)
 - Commercial ADA for Legacy Instrument codes
 - Shared access to memory permits dynamic allocation to processor
- Replacement for angler (Louis Nguyen)
- Tesla GPU server processing (Yongxiang Hu)
- Cloud Modeling System (Kuan-man Xu) (IBM)
- Global Climate Modeling System (David Consodine) (IBM)
- RAQMS Modeling System (Jay Al Saadi) (SGI)
- Legacy systems
 - Continue SGI3800 as interim fix for legacy CERES codes (Norm Loeb)
 - Mac G5 cluster

• Storage

- IBM Read-only access to ASDC Storage via fiberchannel
- Read/Write access to User directed IBM storage via fiberchannel
 - 330TB (usable) RAID-5 in 2 Chassis
 - Allocated by Instrument/subsystem team
 - Expansion capacity to 10 PetaBytes
- Legacy Systems
 - Continue usage of existing legacy Apple SAN systems

• Communications

- Internal Fiberchannel dual 4GB FC for IBM
- External NFS, Samba, AFP via GigE from multiple file servers
- Continued NFS, Samba, AFP via GigE for legacy storage systems

• Support

- 24x7x365 for 5 year

Production Use Model



- Minimize idle time on processors
 - Guarantee minimal levels of processor availability to projects based on funding
 - Prioritize production requests
- Encourage homogeneity among supported science team processing environment
 - Suse version of Linux
- Recognize that one size does not fit all
 - Big endian and little endian processors
 - Small number of big files vs large number of small files
 - Hardware not suitable for all needs

